

Sparse Modeling of Intrinsic Correspondences

J. Pokrass¹, A. M. Bronstein¹, M. M. Bronstein²,
P. Sprechmann³, and G. Sapiro¹

¹School of Electrical Engineering, Tel Aviv University

²Faculty of Informatics, Università della Svizzera Italiana

³Department of Electrical and Computer Engineering, Duke University

October 1, 2012

Abstract

We present a novel sparse modeling approach to non-rigid shape matching using only the ability to detect repeatable regions. As the input to our algorithm, we are given only two sets of regions in two shapes; no descriptors are provided so the correspondence between the regions is not known, nor we know how many regions correspond in the two shapes. We show that even with such scarce information, it is possible to establish very accurate correspondence between the shapes by using methods from the field of sparse modeling, being this, the first non-trivial use of sparse models in shape correspondence. We formulate the problem of *permuted sparse coding*, in which we solve simultaneously for an unknown permutation ordering the regions on two shapes and for an unknown correspondence in functional representation. We also propose a robust variant capable of handling incomplete matches. Numerically, the problem is solved efficiently by alternating the solution of a linear assignment and a sparse coding problem. The proposed methods are evaluated qualitatively and quantitatively on standard benchmarks containing both synthetic and scanned objects.

1 Introduction

Matching of deformable shapes is a notoriously difficult problem playing an important role in many applications [KZHCO10]. Unlike rigid matching where the correspondence can be parametrized by a small number of parameters (rotation and translation of one shape w.r.t. the other [CM91, BM92]), non-rigid matching typically uses point-wise representation of correspondence, which results in the number of degrees of freedom growing exponentially with the number of matched points.

Non-rigid correspondence methods try to find correspondence by minimizing some structure distortion. The structures can be point-wise (local descriptors

[ZBVH09, SOG09, GBAL09, ASC11]), pair-wise (distances [EK01, MS05, BBK06, BBK⁺10]), or higher order [ZWW⁺10].

In order to make the matching problem computationally feasible, it is crucial to reduce the size of the search space [TBW⁺11]. Most methods use a combination of point- and pair-wise structure matching in order to achieve this, and typically consist of three main components: feature detection, feature description, and regularization. Given two shapes, a *feature detector* allows to find a set of landmarks (points or regions) that are repeatable, i.e., appear (possibly with some inaccuracy) on both shapes. *Feature descriptor* then assigns to each feature a vector capturing some local geometric properties of the shape; very often, the two processes are combined into a single one. Using the descriptors, landmarks on two shapes can be matched (it has been shown [OMMG10] that under some conditions, correct landmark matching fully determines the intrinsic correspondence between the shapes). Such a matching reduces the search space size to points with similar descriptors. However, since the matching uses only local information, such correspondence can be noisy, and some kind of *regularization* based on higher-order information is needed to rule out bad or inconsistent correspondences. This information is also used to establish the correspondence between the rest of the points on the shapes. Often, the process is applied hierarchically, restricting the candidate matches to points in the proximity of the landmarks [SY12].

Computer graphics and geometry processing literature contains a plethora of approaches for each of the aforementioned components. Feature detection methods try to locate stable points or regions [DMAMS10, LBB11] that are invariant under isometric deformations and robust to noise. Popular feature descriptors include the heat kernel signature (HKS) [SOG09, GBAL09], wave kernel signature (WKS) [ASC11], global point signature (GPS) [Rus07] or methods adopted from the domain of image analysis [ZBVH09]. As regularization, pair-wise structures such as geodesic [MS05, BBK06] or diffusion distances [BBK⁺10] and higher-order structures [ZWW⁺10] have been used.

Alternatively, there have been several attempts to represent correspondences with a small set of parameters. Elad and Kimmel [EK01] used multidimensional scaling (MDS)-type methods to embed the intrinsic structure of the shapes into a low-dimensional Euclidean space, posing the problem of non-rigid matching as a problem of rigid matching of the corresponding embeddings (“canonical forms”). Mateus et al. [MHK⁺08] used spectral embeddings instead of MDS. Lipman and Funkhouser [LF09] embedded the shapes into a disk by means of conformal maps and represented the correspondence as a Möbius transformation.

More recently, Ovsjanikov et al. [OBCS⁺12] introduced the functional representation of correspondences, allowing to perform a “calculus” of correspondences. In this approach, correspondence is modeled as a correspondence between functions on two shapes rather than points, and can be compactly represented in the Laplace-Beltrami eigenbasis [Lév06] as a matrix of coefficients of decomposition of the basis functions of the first shape in the basis of the second one. In this paper, we will be relying upon this latter representation.

1.1 Main contribution

The main practical contribution of this paper is an approach for finding dense intrinsic correspondence between near-isometric shapes with very little known information: we only assume to be able to detect regions in two shapes in a repeatable enough way (i.e., that at least some regions in one shape correspond accurately enough to some other regions in another shape). No region descriptors are given, so the correspondence of the regions is unknown. The assumption of near-isometry assures that in the functional representation of [OBCS⁺12], the unknown correspondence can be represented as a sparse matrix. The assumption of repeatable regions implies that there exists some unknown permutation that orders the regions according to their correspondence.

We formulate the problem of *permuted sparse coding*, in which we simultaneously look for the permutation and the correspondence, thereby introducing the very successful area of sparse modeling into efficient and state-of-the-art shape correspondence. We note that with the permutation fixed, our problem becomes the standard sparse coding problem; having the correspondence fixed, the problem becomes a linear assignment. This allows efficient numerical solution by alternating the two aforementioned problems and employing efficient solvers that exist for both.

Our method relies on a pretty common assumption that the shapes are nearly-isometric (though our experimental results show our approach still works even when departing from this assumption), and out of all methods we are aware of, it uses perhaps the scarcest amount of data to establish dense correspondence between the shapes. For example, standard region detectors with high repeatability such as [LBB11] are sufficient.

Compared to recent techniques for region-wise shape matching (see, e.g., [GF09, VKTS⁺11, HKG11, PBB11]), our approach has several important practical advantages: First, we do not use any feature descriptor. Second, most region-wise correspondence approaches require an additional step of extending the correspondence between matched regions to the rest of the points.

The rest of the paper is organized as follows. In Section 2, we overview the functional representation of correspondences, allowing to work with correspondences as algebraic structures, and state the main notions in sparse modeling. In Section 3, we formulate our problem of permuted sparse coding for establishing correspondence from a set of repeatable regions given in unknown order. We then extend the problem to the general setting where the region detection process is not perfectly repeatable. In Section 4, we describe the numerical optimization used to solve our permuted sparse coding problem. Experimental results are shown in Section 5. Finally, Section 6 discusses the limitations and possible extensions of the proposed framework and concludes the paper.




Figure 1: Near isometric shape correspondence as a sparse modeling problem (see details in text): Indicator functions of repeatable regions on two shapes are detected and represented as matrices of coefficients \mathbf{A} and \mathbf{B} in the corresponding orthonormal harmonic bases Φ and Ψ . When the regions are brought into correspondence, the point-to-point correspondence between the shapes can be encoded by an approximately diagonal matrix \mathbf{C} . In the proposed procedure termed as *permuted sparse coding*, we solve $\Pi \mathbf{B} = \mathbf{A} \mathbf{C} + \mathbf{O}$ simultaneously for an approximately diagonal \mathbf{C} and the permutation Π bringing the indicator functions into correspondence. To cope with imperfectly matching regions, we relax the surjectivity of the permutation and absorb the mismatches into a row-wise sparse outlier matrix \mathbf{O} . For visualization purposes, the coloring of the regions is consistent as after the application of the permutation. Correspondence is shown between a synthetic TOSCA and scanned SCAPE shape.




Figure 2: Top row: representation of different maps between the two shapes using the matrix \mathbf{C} . Shown left-to-right are the ideal correspondence, a symmetric correspondence, and a random correspondence. Bottom row: representation of the same correspondences as point-to-point maps. Note that the farther is the correspondence from an isometry, the less diagonal is the matrix \mathbf{C} .

2 Background

2.1 Functional representation of correspondences

The direct representation of correspondences as maps between two non-Euclidean spaces limits the range of tools that can be employed for correspondence com-



Figure 3: Dense point-to-point correspondences obtained between the left TOSCA human shape and its approximate isometries. Corresponding points are marked with consistent colors. The average correspondence distortion is depicted in units of the reference shape diameter. The highest distortions are obtained on the non-isometric joints, but do not exceed 6% of the diameter.

putation due to the lack of an algebraic structure. In this paper, we rely on the functional representation of correspondences introduced in [OBCS⁺12], which overcomes this limitation. In what follows, we briefly review the main idea of such functional representations.

Let X and Y be two shapes, modeled as compact smooth Riemannian manifolds, related by a bijective correspondence $t : X \rightarrow Y$. Then, for any real function $f : X \rightarrow \mathbb{R}$, we can construct a corresponding function $g : Y \rightarrow \mathbb{R}$ as $g = f \circ t^{-1}$. The correspondence t uniquely defines a mapping between two function spaces $T : \mathcal{F}(X, \mathbb{R}) \rightarrow \mathcal{F}(Y, \mathbb{R})$, where $\mathcal{F}(X, \mathbb{R})$ denotes the space of real functions on X . Such a representation is linear, since for every pair of functions f_1, f_2 and scalars α_1, α_2 ,

$$\begin{aligned} T(\alpha_1 f_1 + \alpha_2 f_2) &= (\alpha_1 f_1 + \alpha_2 f_2) \circ t^{-1} \\ &= \alpha_1 f_1 \circ t^{-1} + \alpha_2 f_2 \circ t^{-1} \\ &= \alpha_1 T(f_1) + \alpha_2 T(f_2). \end{aligned} \quad (1)$$

Assuming that X is equipped with a basis $\{\phi_i\}_{i \geq 1}$, any $f : X \rightarrow \mathbb{R}$ can be represented as

$$f = \sum_{i \geq 1} a_i \phi_i \quad (2)$$

with the a_i being some representation coefficients (in case of an orthonormal basis, $a_i = \langle f, \phi_i \rangle$; in the general case, the coefficients are found by projecting the function f on the bi-orthonormal basis). Due to the linearity of T ,

$$T(f) = T\left(\sum_{i \geq 1} a_i \phi_i\right) = \sum_{i \geq 1} a_i T(\phi_i) \quad (3)$$

If the shape Y is further equipped with a basis $\{\psi_j\}_{j \geq 1}$, then for every i there

exists coefficients c_{ij} such that

$$T(\phi_i) = \sum_{j \geq 1} c_{ij} \psi_j, \quad (4)$$

and we can write

$$T(f) = \sum_{i,j \geq 1} a_i c_{ij} \psi_j. \quad (5)$$

Let us now assume finite sampling of X and Y , with m samples (for simplicity, we assume that the shapes are sampled at the same number of samples m). The extension to the case with a different number of samples is straightforward). The bases are represented as the $m \times n$ matrices Φ and Ψ containing, respectively, n discretized functions ϕ_i and ψ_j as the columns. The functions f and $g = T(f)$ can now be represented as n -dimensional vectors $\mathbf{f} = \Phi \mathbf{a}$ and $\mathbf{g} = \Psi \mathbf{b}$ with the coefficients \mathbf{a} and \mathbf{b} . Using this notation, Equation (5) can be rewritten as

$$\Psi \mathbf{b} = T(\Phi \mathbf{a}) = \Psi \mathbf{C}^T \mathbf{a}; \quad (6)$$

since Ψ is invertible, this simply means that

$$\mathbf{b}^T = \mathbf{a}^T \mathbf{C}. \quad (7)$$

Thus, the $n \times n$ matrix \mathbf{C} fully encodes the linear map T between the functional spaces, and contains the coordinates in the basis Ψ of the mapped elements of the basis Φ .

2.2 Point-to-point correspondence

Point-to-point correspondences assume that each point i on X corresponds to some point j on Y . In functional representation, this is equivalent to having \mathbf{C} that makes each row of $\Psi \mathbf{C}^T$ coincide with some row of Φ [OBCS⁺12]. In applications requiring point-to-point correspondence, given some \mathbf{C} , it can be converted into a point-to-point correspondence by thinking of Φ and Ψ as n -dimensional points clouds, and orthogonal matrix \mathbf{C} as a rigid alignment transformation between them. This procedure is equivalent to iterative closest point (ICP) in n dimensions [OBCS⁺12], initialized with the given \mathbf{C}_0 : first, for each row i of $\Psi \mathbf{C}_0^T$, find the closest row j_i^* in Φ (this operation can be performed efficiently using approximate nearest neighbor algorithms). Then, find orthonormal \mathbf{C} minimizing $\sum_i \|\Phi_{j_i^*} - \Psi \mathbf{C}^T\|_2$ and set $\mathbf{C}_0 = \mathbf{C}$. This operation is repeated until convergence and can be performed efficiently over all the vertexes of X and Y using approximate nearest neighbor algorithms.

2.3 Sparse modeling

One of the main tools that will be used in this paper are *sparse models*. In what follows, we give a very brief overview of this vast field, and refer the reader to [Ela10] for a comprehensive treatise. The central assertion of sparse modeling

is that many families of signals (and later operations as here introduced) can be represented as a sparse linear combination in an appropriate domain, usually referred to as the *dictionary*. This can be written as $\mathbf{x} \approx \mathbf{D}\mathbf{z}$, where \mathbf{x} denotes the signal, \mathbf{D} the dictionary, and \mathbf{z} the sparse vector of representation coefficients. The dictionary is often selected to be *overcomplete*, i.e., with more columns than rows.

Finding the representation of a signal \mathbf{x} in a given dictionary \mathbf{D} is usually referred to as sparse representation *pursuit* or *sparse coding*. Among the variety of pursuit methods, we will focus on the so-called Lasso formulation [Tib96] that finds \mathbf{z} as the solution to the unconstrained convex program

$$\min_{\mathbf{z}} \|\mathbf{x} - \mathbf{D}\mathbf{z}\|_2^2 + \lambda \|\mathbf{z}\|_1. \quad (8)$$

The first term is the data fitting term, while the second term involving ℓ_1 norm, $\|\mathbf{z}\|_1 = |z_1| + \dots + |z_n|$, promotes a sparse solution; the parameter λ controls the relative importance of the latter. Proximal splitting methods [Nes07] are among the most efficient and most frequently used numerical tools to solve problem (8); in Section 4, we present a variant of the proximal splitting algorithms for the solution of the pursuit problem arising in shape correspondence as detailed in the sequel.

In some cases, signals not admitting the simplistic model of element-wise sparsity still manifest more intricate types of *structured* sparsity. In structured sparse models, the non-zero elements of \mathbf{z} come in groups or, more generally, in hierarchies of groups. A common class of structured pursuit problems can be formulated as convex programs of the form

$$\min_{\mathbf{z}} \|\mathbf{x} - \mathbf{D}\mathbf{z}\|_2^2 + \lambda \|\mathbf{z}\|_{1,2}, \quad (9)$$

where the $\ell_{1,2}$ norm, $\|\mathbf{z}\|_{1,2} = \|\mathbf{z}_1\|_2 + \dots + \|\mathbf{z}_k\|_2$, assumes that the vector \mathbf{z} is decomposed into k non-overlapping sub-vectors \mathbf{z}_i , and promotes group-wise sparse solutions (i.e., the solution will have a small number of groups with non-zero coefficients, but the sub-vectors representing each such non-zero group will be dense).

While structured sparse models enforce group structure of each representation vector independently, it is often useful to consider the structure shared by multiple vectors. *Collaborative* sparse models operate on data matrices \mathbf{X} , in which each column corresponds to a data vector, and assert that the patterns of non-zero coefficients are shared across the corresponding representation vectors, \mathbf{Z} . This is achieved by solving a pursuit problem of the form

$$\min_{\mathbf{Z}} \|\mathbf{X} - \mathbf{D}\mathbf{Z}\|_F^2 + \lambda \|\mathbf{Z}\|_{2,1}, \quad (10)$$

where the first term involving the Frobenius norm serves as the data fitting term, and the second term with the $\ell_{2,1}$ norm promotes row-wise sparsity of the solution. The $\ell_{2,1}$ norm is defined as $\|\mathbf{Z}\|_{2,1} = \|\mathbf{z}_1^T\|_2 + \dots + \|\mathbf{z}_m^T\|_2$, where \mathbf{z}_i^T denotes the i -th row of \mathbf{Z} (note the difference from the $\ell_{1,2}$ column-wise counterpart!).



Figure 4: Dense point-to-point correspondences obtained between the left SCAPE human shape and various other poses. Corresponding points are marked with consistent colors.




Figure 5: First row: point-to-point correspondences obtained between different non-isometric shapes: male and female (left); two strongly non-isometric deformations of the dog shape from the TOSCA set (middle); TOSCA and SCAPE human shapes (right). Second row: Point-to-point correspondences obtained between SHREC shapes undergoing nearly isometric deformations and (from left to right) spike noise, Gaussian noise, and topological noise in the form of large and small holes.

In this paper, we use formulate the shape correspondence problem using a sparse model, and use sparse modeling tools to efficiently solve it.

3 Sparse modeling of correspondences

In case the shapes X and Y are isometric and the corresponding Laplace-Beltrami operators have simple spectra (no eigenvalues with multiplicity greater than one), the harmonic bases (Laplacian eigenfunctions) have a compatible be-

havior, $\psi_i = T(\phi_i)$ such that $c_{ij} = \pm \delta_{ij}$. Choosing the discretized eigenfunctions of the Laplace-Beltrami operator as Φ and Ψ causes every low-distortion correspondence being represented by a nearly diagonal, and therefore very sparse, matrix \mathbf{C} .

In practice, due to lack of perfect isometry and numerical inaccuracies, the diagonal structure of \mathbf{C} is manifested for the first eigenfunctions corresponding to the small eigenvalues (low frequencies), and is gradually lost with the increase of the frequency (Figure 2). However, a correspondence with low metric distortion will usually be represented by a sparse \mathbf{C} . We use this property to formulate the computation of correspondences in terms of a sparse representation pursuit problem.

Let us assume to have some *region* (or *feature*) *detection* process that given a shape X produces a collection of functions $\{f_i : X \rightarrow \mathbb{R}\}$ based on the intrinsic properties of the shape only. For example, the f_i 's can be indicator functions of the maximally stable components (regions) of the shape [LBB11]. Since the process is intrinsic, given a nearly isometric deformation Y or X , it should produce a collection of similar functions $\{g_j : Y \rightarrow \mathbb{R}\}$.

To simplify the presentation, let us assume that the process is perfectly *repeatable* in the sense that it finds q functions on X and Y , such that for every f_i there exists a $g_j = f_i \circ t$ related by the unknown correspondence t . We stress that the ordering of the f_i 's and g_j 's is *unknown*, i.e., we do not know to which g_j in Y a f_i in X correspond. This ordering can be expressed by an unknown $q \times q$ permutation matrix Π (in Section 3.2, we consider the more general case when the number of functions detected on X and Y can be different, i.e., Π is non-square).

Representing the functions in the bases on each shape, we have $\mathbf{f}_i = \Phi \mathbf{a}_i$ and $\mathbf{g}_j = \Psi \mathbf{b}_j$. Since each pair of corresponding \mathbf{f}_i and \mathbf{g}_j shall satisfy (7), we can write

$$\Pi \mathbf{B} = \mathbf{A} \mathbf{C}, \quad (11)$$

where \mathbf{A} and \mathbf{B} are the $q \times n$ matrices containing, respectively, \mathbf{a}_i^T and \mathbf{b}_j^T as the rows, and $\pi_{ij} = 1$ if \mathbf{a}_i corresponds to \mathbf{b}_j and zero otherwise.

3.1 Permuted sparse coding

Note that in relation (11), both Π and \mathbf{C} are unknown, and solving for them is a highly ill-posed problem. However, by recalling that the correspondence we are looking for should be represented by a nearly-diagonal \mathbf{C} , we formulate the following problem

$$\min_{\mathbf{C}, \Pi} \frac{1}{2} \|\Pi \mathbf{B} - \mathbf{A} \mathbf{C}\|_F^2 + \lambda \|\mathbf{W} \odot \mathbf{C}\|_1, \quad (12)$$

where the minimum is sought over $n \times n$ matrices \mathbf{C} (capturing the correspondence t between the shapes in the functional representation) and $q \times q$ permutations Π (capturing the correspondence between the detected regions on the

shapes). The first term containing the Frobenius norm can be interpreted as the data term, while the second term, containing the weighted ℓ_1 norm promotes a sparse \mathbf{C} ; \odot denotes element-wise multiplication, and the non-negative parameter λ determines the relative importance of the penalty. Small weights w_{ij} in \mathbf{W} are assigned close to the diagonal, while larger weights are selected for the off-diagonal elements. This promotes diagonal solutions.

The solution of (12) can be obtained using alternating minimization iterating over \mathbf{C} with fixed $\mathbf{\Pi}$, and $\mathbf{\Pi}$ with fixed \mathbf{C} . Note that with fixed $\mathbf{\Pi}$, we can denote $\mathbf{B}' = \mathbf{\Pi B}$ and reduce problem (12) to

$$\min_{\mathbf{C}} \frac{1}{2} \|\mathbf{B}' - \mathbf{AC}\|_F^2 + \lambda \|\mathbf{W} \odot \mathbf{C}\|_1, \quad (13)$$

which resembles the Lasso problem frequently employed for the pursuit of sparse representations. On the other hand, when \mathbf{C} is fixed, we set $\mathbf{A}' = \mathbf{AC}$, reducing the optimization objective to

$$\begin{aligned} \|\mathbf{\Pi B} - \mathbf{A}'\|_F^2 = \\ \text{tr}(\mathbf{B}^T \mathbf{\Pi}^T \mathbf{\Pi B}) - 2\text{tr}(\mathbf{B}^T \mathbf{\Pi}^T \mathbf{A}') + \text{tr}(\mathbf{A}'^T \mathbf{A}'). \end{aligned} \quad (14)$$

Since $\mathbf{\Pi}$ is a permutation matrix, $\mathbf{\Pi}^T \mathbf{\Pi} = \mathbf{I}$, and the only non-constant term remaining in the objective is the second linear term. Problem (12) thus becomes

$$\max_{\mathbf{\Pi}} \text{tr}(\mathbf{\Pi}^T \mathbf{E}), \quad (15)$$

where $\mathbf{E} = \mathbf{A}' \mathbf{B}^T = \mathbf{ACB}^T$ and the maximization is performed over permutation matrices. To make it practically solvable, we allow $\mathbf{\Pi}$ to be a double-stochastic matrix, which yields the following linear assignment problem:

$$\max_{\mathbf{\Pi} \geq \mathbf{0}} \text{vec}(\mathbf{E})^T \text{vec}(\mathbf{\Pi}) \quad \text{s.t.} \quad \begin{cases} \mathbf{\Pi} \mathbf{1} = \mathbf{1} \\ \mathbf{\Pi}^T \mathbf{1} = \mathbf{1}. \end{cases} \quad (16)$$

We refer to problem (12) as to *permuted sparse coding*, and propose to solve it by alternating the solution of the standard sparse coding problem (13) and the solution of the linear assignment problem (16). The sparsity constraint has a regularization effect on this, otherwise extremely ill-posed, problem, and the following strong property holds:

Proposition 1. *The process alternating subproblems (13) and (16) converges to a global minimizer of the permuted sparse coding problem (12).*

Due to lack of space, we provide the proof in the Appendix. This result means, among other, that despite the relaxation of the permutation matrix to a double-stochastic matrix in the assignment subproblem (16), we are actually recovering a true permutation matrix. This follows from the total unimodularity of the constraints in (16).

3.2 Robust permuted sparse coding

So far, we have assumed the existence of a bijective, albeit unknown, correspondence between the f_i 's and the g_j 's. In practice, the process detecting these functions (e.g., stable regions) is often not perfectly repeatable. In what follows, we will make a more realistic assumption that q functions f_i are detected on X , and r functions g_j detected on Y (without loss of generality, $q \leq r$), such that some f_i 's have no counterpart g_j , and vice versa. This partial correspondence can be described by a $q \times r$ partial permutation matrix $\mathbf{\Pi}$ in which now some columns and rows may vanish.

Let us assume that $s \leq q$ f_i 's have corresponding g_j 's. This means that there is no correspondence between $r - s$ rows of \mathbf{B} and $q - s$ rows of \mathbf{A} , and the relation $\mathbf{\Pi B} \approx \mathbf{A C}$ holds only for an unknown subset of its rows. The mismatched rows of \mathbf{B} can be ignored by letting some columns of $\mathbf{\Pi}$ vanish, meaning that the correspondence is no more surjective. This can be achieved by relaxing the equality constraint $\mathbf{\Pi}^T \mathbf{1} = \mathbf{1}$ in (16) replacing it with $\mathbf{\Pi}^T \mathbf{1} \leq \mathbf{1}$. However, dropping injectivity as well and relaxing $\mathbf{\Pi 1} = \mathbf{1}$ to $\mathbf{\Pi 1} \leq \mathbf{1}$ would result in the trivial solution $\mathbf{\Pi} = \mathbf{0}$. To overcome this difficulty, we demand every row of \mathbf{A} to have a matching row in \mathbf{B} , and absorb the $r - s$ mismatches in a row-sparse $q \times n$ outlier matrix \mathbf{O} that we add to the data term of (12). This results in the following problem

$$\min_{\mathbf{C}, \mathbf{O}, \mathbf{\Pi}} \frac{1}{2} \|\mathbf{\Pi B} - \mathbf{A C} - \mathbf{O}\|_F^2 + \lambda \|\mathbf{W} \odot \mathbf{C}\|_1 + \mu \|\mathbf{O}\|_{2,1}, \quad (17)$$

which we refer to as *robust permuted sparse coding*. The last term involves the $\ell_{2,1}$ norm

$$\|\mathbf{O}\|_{2,1} = \sum_{i=1}^r \|\mathbf{o}_i^T\|_2, \quad (18)$$

which can be thought of as the ℓ_1 norm of the vector of the ℓ_2 norms of the rows \mathbf{o}_i^T of \mathbf{O} . The $\ell_{2,1}$ norm promotes row-wise sparsity, allowing to absorb the errors in the data term corresponding to the rows of \mathbf{A} having no corresponding rows in \mathbf{B} ; the parameter $\mu \geq 0$ controls the amount of regularization. The $q \times r$ matrix $\mathbf{\Pi}$ is searched over all injective correspondences.

As before, problem (17) is split into two sub-problems, one with the fixed permutation $\mathbf{\Pi}$,

$$\min_{\mathbf{C}, \mathbf{O}} \frac{1}{2} \|\mathbf{B}' - \mathbf{A C} - \mathbf{O}\|_F^2 + \lambda \|\mathbf{W} \odot \mathbf{C}\|_1 + \mu \|\mathbf{O}\|_{2,1}, \quad (19)$$

with $\mathbf{B}' = \mathbf{\Pi B}$, and the other one with the fixed \mathbf{C} ,

$$\max_{\mathbf{\Pi} \geq \mathbf{0}} \text{vec}(\mathbf{E})^T \text{vec}(\mathbf{\Pi}) \quad \text{s.t.} \quad \begin{cases} \mathbf{\Pi 1} = \mathbf{1} \\ \mathbf{\Pi}^T \mathbf{1} \leq \mathbf{1}, \end{cases} \quad (20)$$

with $\mathbf{E} = (\mathbf{A C}) \mathbf{B}^T$. Note that an injective correspondence is relaxed into a row-wise stochastic and column-wise sub-stochastic matrix $\mathbf{\Pi}$. Proposition 1 simply extends to the robust formulation as well.

4 Numerical solution

The solution of the robust permuted sparse coding problem (17) is reduced to alternating two relatively standard optimization problems, and there exist many readily available efficient numerical tools to solve them. For the sake of completeness, we provide a concise description of the involved numerics.

Problem (20), being a simple linear assignment problem, is solved using the Hungarian algorithm. As an alternative, linear programming can be employed. To reduce the search space size, we disallow certain impossible permutations such as those relating regions with very distinct sizes.

In order to solve (19), we use the family of forward-backward splitting algorithms [Nes07] designed for solving unconstrained optimization problems in which the cost function can be split into the sum of two terms,

$$\min_{\mathbf{x}} h_1(\mathbf{x}) + h_2(\mathbf{x}), \quad (21)$$

one, h_1 , convex and differentiable with an α -Lipschitz continuous gradient and another, h_2 , convex extended real valued and possibly non-smooth. Clearly, problem (19) falls in this category.

The forward-backward splitting method with fixed constant step defines a series of iterates, $\{\mathbf{x}^k\}_k$,

$$\mathbf{x}^{k+1} = \mathbf{P}_{\alpha h_2} \left(\mathbf{x}^k - \frac{1}{\alpha} \nabla h_1(\mathbf{x}^k) \right), \quad (22)$$

where

$$\mathbf{P}_{\alpha h_2}(\mathbf{x}) = \arg \min_{\mathbf{u}} \|\mathbf{u} - \mathbf{x}\|_2^2 + \alpha h_2(\mathbf{u}) \quad (23)$$

denotes the proximal operator of h_2 . Many alternatives have been studied in the literature to improve the convergence rate of forward-backward splitting algorithms [BT09, Nes07]. Accelerated versions reach quadratic convergence rates (the best possible for the class of first order methods). The discussion of these methods is beyond of the scope of this paper.

In our case, the objective comprises a quadratic function $h_1 = \|\mathbf{B}' - \mathbf{A}\mathbf{C} - \mathbf{O}\|_F^2$ and the non-smooth function $h_2 = \lambda \|\mathbf{W} \odot \mathbf{C}\|_1 + \mu \|\mathbf{O}\|_{2,1}$. The proximal operator splits into two operators, one in \mathbf{C} and another one in \mathbf{O} , both having closed forms. The proximal operator corresponding to the ℓ_1 norm term is given by the weighted soft threshold function

$$\mathbf{P}_1(\mathbf{C}) = \max \left\{ |\mathbf{C}| - \frac{\lambda}{\alpha} \mathbf{W} \right\} \odot \text{sign}(\mathbf{C}), \quad (24)$$

where the absolute value and the sign functions are applied element-wise. The i -th row of the proximal operator corresponding to the $\ell_{2,1}$ norm term is given by

$$(\mathbf{P}_2(\mathbf{O}))_i = \max \left\{ \|\mathbf{o}_i^T\|_2 - \frac{\mu}{\alpha} \right\} \frac{\mathbf{o}_i^T}{\|\mathbf{o}_i^T\|_2}. \quad (25)$$

input : Data \mathbf{B}', \mathbf{A} ; parameters λ, μ ; step size α .
output: Sparse matrix \mathbf{O} and row-wise sparse outlier matrix \mathbf{O}
Initialize $\mathbf{O}^0 = \mathbf{B}'$ and $\mathbf{C}^0 = \mathbf{0}$.
for $k=1, 2, \dots$, *until convergence* **do**
 $\mathbf{C}^{k+1} = \mathbf{P}_1 \left((\mathbf{I} - \frac{1}{\alpha} \mathbf{A}^T \mathbf{A}) \mathbf{C}^k - \frac{1}{\alpha} \mathbf{A}^T (\mathbf{O}^k - \mathbf{B}') \right)$
 $\mathbf{O}^{k+1} = \mathbf{P}_2 \left((1 - \frac{1}{\alpha}) \mathbf{O}^k - \frac{1}{\alpha} (\mathbf{A} \mathbf{C}^k - \mathbf{B}') \right)$
end

Algorithm 1: Forward-backward splitting method for the solution of (19).

The gradient of the quadratic data term with respect to \mathbf{C} and \mathbf{O} is given straightforwardly by

$$\begin{aligned}
\nabla_{\mathbf{C}} h_1 &= \mathbf{A}^T \mathbf{A} \mathbf{C} + \mathbf{A}^T \mathbf{O} - \mathbf{A}^T \mathbf{B}' \\
\nabla_{\mathbf{O}} h_1 &= \mathbf{O} + \mathbf{A} \mathbf{C} - \mathbf{B}'.
\end{aligned} \tag{26}$$

The Lipschitz constant of the gradient determining the step size is bounded by the maximum eigenvalue

$$\alpha \leq \lambda_{\max} \begin{pmatrix} \mathbf{A}^T \mathbf{A} & \mathbf{A}^T \\ \mathbf{I} & \mathbf{A} \end{pmatrix}. \tag{27}$$

Plugging the above expressions together into (22) yields the forward-backward splitting optimization summarized in Algorithm 1.

5 Experimental results

In order to evaluate our approach, we performed several experiments on data from the TOSCA [BBK08], SHREC'11 [EB] and SCAPE [ASK⁺05] datasets. The TOSCA set contains high-quality (10K-50K vertices) synthetic triangular meshes of humans and animals in different poses with known ground truth correspondences between them. SHREC'11 contains meshes from the TOSCA set undergoing simulated transformations. The SCAPE set contains high-resolution (12K vertices) scans of a real human in different poses.

For each pair of shapes we calculated the MSERs using 6-10 eigenfunctions and selected regions with areas of at least 5-10% of the total shape area, resulting in about 5 – 15 detected regions (see Figure 1). These parameters were selected empirically for our data sets.

The segments of each shape were projected onto 20 eigenfunctions and the corresponding \mathbf{C} matrix was calculated by solving the sparse coding subproblem (19) using an accelerated variant of the method described in Section 4. The linear assignment subproblem (16) was solved using the Hungarian method [Kuh55]. We initialized the permutation matrix with $\mathbf{\Pi} = \frac{1}{q} \mathbf{1} \mathbf{1}^T$, and the correspondence matrix with $\mathbf{C} = \mathbf{0}$. We observed a rapid convergence of the alternating minimization procedure in one or two iterations (see Figure 6 where for visualization purposed, $\mathbf{\Pi}$ was initialized to identity). We found that the




Figure 6: Outer iterations of robust permuted sparse coding alternating the solution of the sparse representation pursuit problem (19) with the linear assignment problem (20). Three iterations, shown left-to-right, are required to achieve convergence. Depicted are the permutation matrix $\mathbf{\Pi}$ (first row), the correspondence matrix \mathbf{C} (second row), and the outlier matrix \mathbf{O} (last row). The resulting point-to-point correspondence and the correspondence matrix \mathbf{C} refined using the ICP as described in Section 2.2 are shown in the rightmost column.

method consistently converged to the same solution regardless of the initialization. Finally, after convergence of the alternating minimization, the resulting \mathbf{C} was refined using the method described in Section 2.2.

The robustness of the method is demonstrated in Figures 3–5; correct correspondences are computed even when the shapes undergo non-isometric deformations and are contaminated by geometric or topological noise. Figure 7 shows a quantitative evaluation and comparison of our algorithm to other correspondence algorithms on the SCAPE data set. The evaluation was performed using the code and data from [KLF11]. Our method outperforms existing methods while making less assumption and working only with intrinsic information.

The code used in the experiments was implemented in Matlab with parts written in C. The approximate nearest neighbor search in the ICP refinement step was accelerated using the FLANN library. The experiments were run on a 2.4GHz Intel Xeon CPU. End-to-end execution time varied from 10 to 50 seconds, with the detailed breakdown summarized in Table 1.

6 Discussion and Conclusion

In this paper, we posed the problem of finding intrinsic correspondence between near-isometric deformable shapes as a problem of sparse modeling. Given only two set of regions in the two shapes with unknown correspondence, we are




Figure 7: Quantitative evaluation of the proposed shape correspondence algorithm and its comparison to other correspondence algorithms on the SCAPE shapes using the evaluation protocol from [KLF11].

able to infer a dense correspondence between the shapes from two assumptions: that at least some of the regions in the two sets are corresponding; and that the shapes are nearly-isometric. The latter assumption implies that in functional representation in harmonic bases the unknown correspondence between the shapes is modeled as a sparse nearly-diagonal matrix; the former assumption implies that there exists an unknown permutation that reorders the regions in corresponding order. To find both the permutation and the correspondence, we formulate the novel permuted sparse coding problem and propose its efficient solution. An additional sparse coding term addressing outliers is added to the model for handling partial matching and formulated as the robust permuted sparse coding.

To the best of our knowledge, among other dense correspondence techniques, our method relies on the smallest amount of information (the ability to find some repeatable regions) and quite generic assumption (near-isometric shapes). In particular, it allows us to use only a region detector without a feature descriptor to find a high-quality correspondence between two shapes.

We note that, as in [OBCS⁺12], we explicitly assume that the shapes are nearly isometric, and that their Laplacians have simple spectrum. This assumption assures that the Laplacian eigenbases Φ and Ψ have a compatible behavior, and as a result \mathbf{C} has a nearly-diagonal structure. If we try to relax the restriction on multiplicity, \mathbf{C} will still be sparse, but with unknown sparse structure. We can still use our problem in this setting, imposing a different sparsity

Table 1: Average runtime (in seconds) as a function of the shape size for different stages in the proposed method: Basis – harmonic basis computation; MSER – region detection; Opt. – alternating minimization procedure; Ref. – ICP-based refinement and point-to-point correspondence computation; Tot. – total runtime.

Vertices	Basis	MSER	Opt.	Ref.	Tot.
5K	0.53	0.61	7.80	1.41	10.35
10K	0.99	1.32	7.91	2.70	12.92
20K	2.03	3.58	7.91	5.52	19.04
50K	5.57	14.23	7.85	13.99	41.64

constraint on \mathbf{C} .

Relaxing the assumptions even more, we can depart from the near-isometric model, e.g. considering applications where one wishes to match shapes with roughly similar geometry but very different details (such as a horse and an elephant). In such a generic setting, the Laplacian eigenbases may differ dramatically, and thus \mathbf{C} have a non-sparse structure. It is possible to incorporate the bases Φ and Ψ as variables into our problem, and in addition to finding the permutation Π and correspondence \mathbf{C} find also the bases in which \mathbf{C} will have a diagonal structure. This problem is akin to dictionary learning used in the sparse modeling literature. In future research, we will study such a generalization of our framework in hope to find correspondences between non-isometric shapes. Another possible generalization of our problem is for finding correspondence between collections of shapes [NBCW⁺11, KLF11].

Finally, it worthwhile noting that the novel structured sparse modeling techniques introduced in [SBS12] provide an alternative to the optimization-based pursuit by replacing the iterative proximal algorithm with a learned fixed-complexity feed-forward network. Approaching shape correspondence as a learning problem from this perspective seems a very attractive future research direction.

Appendix A – Proof of Proposition 1

The permuted sparse coding problem

$$\min_{\mathbf{C} \in \mathbb{R}^{q \times q}, \Pi \in \mathbb{P}(q)} \frac{1}{2} \|\Pi \mathbf{B} - \mathbf{A} \mathbf{C}\|_{\text{F}}^2 + \lambda \|\mathbf{W} \odot \mathbf{C}\|, \quad (28)$$

where $\mathbb{P}(q)$ denotes the space of $q \times q$ permutation matrices, is non-convex since the feasible set is non-convex. However, by relaxing $\mathbb{P}(q)$ to the bigger space of $q \times q$ double-stochastic matrices, $\mathbb{S}(q) = \{\Pi \in \mathbb{R}_+^{q \times q} : \Pi \mathbf{1} = \Pi^T \mathbf{1} = \mathbf{1}\} \supset \mathbb{P}(q)$, we obtain the problem

$$\min_{\mathbf{C} \in \mathbb{R}^{q \times q}, \Pi \in \mathbb{S}(q)} \frac{1}{2} \|\Pi \mathbf{B} - \mathbf{A} \mathbf{C}\|_{\text{F}}^2 + \lambda \|\mathbf{W} \odot \mathbf{C}\| \quad (29)$$

that is easily shown to be convex. Fixing one of the variables at a time, the problem can be split into two subproblems: the sparse coding problem

$$\min_{\mathbf{C}} \frac{1}{2} \|\mathbf{B}' - \mathbf{A}\mathbf{C}\|_{\text{F}}^2 + \lambda \|\mathbf{W} \odot \mathbf{C}\|_1 \quad (30)$$

with $\mathbf{B}' = \mathbf{\Pi}\mathbf{B}$, and the linear assignment problem

$$\max_{\mathbf{\Pi} \in \mathbb{S}(q)} \text{vec}(\mathbf{E})^{\text{T}} \text{vec}(\mathbf{\Pi}) \quad (31)$$

with $\mathbf{E} = \mathbf{A}'\mathbf{B}^{\text{T}} = \mathbf{A}\mathbf{C}\mathbf{B}^{\text{T}}$. These two subproblems can be viewed as minimizing the objective function of (29) with respect to two blocks of coordinates, \mathbf{C} and $\mathbf{\Pi}$. The minimization process alternating between the solution of (30) and (31) can be therefore regarded as block-coordinate descent.

We use an instance of Theorem 4.1 in [Tse01] stating that block-coordinate descent is guaranteed to converge to a global minimizer of a non-differentiable convex function. The conditions of the theorem are satisfied by the objective and the constraints of (29). Note that since we do not prove strict convexity of the latter objective, a global minimizer is not necessarily unique.

Let us now have a closer look at the linear assignment problem (31) that can be cast as the linear program

$$\min_{\boldsymbol{\pi} \in \mathbb{R}_+^{q^2}} \mathbf{e}^{\text{T}} \boldsymbol{\pi} \quad \text{s.t.} \quad \mathbf{Q}\boldsymbol{\pi} = \mathbf{1}, \quad (32)$$

where the variable vector $\boldsymbol{\pi} = \text{vec}(\mathbf{\Pi})$ is the column stack of the assignment matrix, the cost vector is given by $\mathbf{e} = -\text{vec}(\mathbf{E})$, and the constraint matrix can be expressed using the Krönecker product notation as

$$\mathbf{Q} = \begin{pmatrix} \mathbf{1}^{\text{T}} \otimes \mathbf{I} \\ \mathbf{I} \otimes \mathbf{1}^{\text{T}} \end{pmatrix}. \quad (33)$$

Here, \mathbf{I} stands for the $q \times q$ identity matrix, and $\mathbf{1}^{\text{T}}$ for the $1 \times q$ vector of ones. In Lemma 1 below, we prove that \mathbf{Q} is totally unimodular, i.e., any of its square submatrices has the determinant of 0 or ± 1 . This property guarantees that the linear program has a solution with integer coordinates. Since $\mathbb{S}(q) \cap \mathbb{Z}^{q \times q} = \mathbb{P}(q)$, this guarantees that the linear assignment problem (31) is actually equivalent to the binary assignment problem

$$\max_{\mathbf{\Pi} \in \mathbb{P}(q)} \text{vec}(\mathbf{E})^{\text{T}} \text{vec}(\mathbf{\Pi}). \quad (34)$$

Combined this result with the convergence of the block-coordinate descent to a global minimizer of (29), we can guarantee that in the obtained minimizer $\mathbf{\Pi}$ actually belongs to $\mathbb{P}(q)$. This implies that the block-coordinate descent converges to a global minimizer of (28).

Lemma 1. *The matrix \mathbf{Q} in (33) is totally unimodular.*

Proof. The matrix \mathbf{Q} can be constructed as the sub-matrix of a bigger matrix

$$\bar{\mathbf{Q}} = \begin{pmatrix} \mathbf{1}^T \\ \mathbf{I} \end{pmatrix} \otimes \begin{pmatrix} \mathbf{I} \\ \mathbf{1}^T \end{pmatrix}. \quad (35)$$

Hence, total unimodularity of $\bar{\mathbf{Q}}$ implies total unimodularity of \mathbf{Q} . Since total unimodularity is preserved by the Krönecker product, it is sufficient to show that each of the two Krönecker factors are totally unimodular. We will limit the discussion to the second factor; very similar arguments apply to the first one.

The matrix

$$\mathbf{R} = \begin{pmatrix} \mathbf{I} \\ \mathbf{1}^T \end{pmatrix} \quad (36)$$

comprises two components: the identity matrix \mathbf{I} and the row vector $\mathbf{1}^T$, both of which are totally unimodular. For each $k \times k$ square submatrix \mathbf{R}' of \mathbf{R} , we distinguish between the following three cases: \mathbf{R}' containing only elements of \mathbf{I} ; \mathbf{R}' containing only elements of $\mathbf{1}^T$; and \mathbf{R}' containing elements of both components. In the two former cases, $\det \mathbf{R}' \in \{0, \pm 1\}$ since \mathbf{I} and $\mathbf{1}^T$ are totally unimodular. In the latter case, the submatrix has the form

$$\mathbf{R}' = \begin{pmatrix} \mathbf{0} & \mathbf{I} \\ \mathbf{1} & \mathbf{1}^T \end{pmatrix}, \quad (37)$$

where the $m \times m$ ($m < k$) identity matrix \mathbf{I} is surrounded by zeros and concatenated to a row of ones on the bottom. For $m < k - 1$, the submatrix contains at least one row of zeros and therefore $\det \mathbf{R}' = 0$. For $m = k - 1$, \mathbf{R}' assumes the form

$$\mathbf{R}' = \begin{pmatrix} \mathbf{0} & \mathbf{I} \\ 1 & \mathbf{1}^T \end{pmatrix}, \quad (38)$$

where \mathbf{I} is the $(k-1) \times (k-1)$ identity matrix, $\mathbf{0}$ is the $(k-1) \times 1$ vector of zeros, and $\mathbf{1}^T$ is the $1 \times (k-1)$ vector of ones. Using the properties of the determinant, we obtain

$$\det \mathbf{R}' = \det(\mathbf{0}\mathbf{1}^T - \mathbf{I}) = (-1)^{k-1}. \quad (39)$$

Hence, \mathbf{R} is totally unimodular. \square

References

- [ASC11] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: a quantum mechanical approach to shape analysis. In *Proc. Workshop on Dynamic Shape Capture and Analysis*, 2011.
- [ASK⁺05] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, , and J. Davis. Scape: shape completion and animation of people. In *Proceedings of the SIGGRAPH Conference*, 2005.

- [BBK06] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS*, 103(5):1168–1172, 2006.
- [BBK08] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. *Numerical geometry of non-rigid shapes*. Springer-Verlag New York Inc, 2008.
- [BBK⁺10] A. M. Bronstein, M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro. A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *IJCV*, 89(2-3):266–286, 2010.
- [BM92] P. J. Besl and N. D. McKay. A method for registration of 3D shapes. *Trans. PAMI*, 14:239–256, 1992.
- [BT09] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Img. Sci.*, 2:183–202, March 2009.
- [CM91] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. In *Proc. Conf. Robotics and Automation*, 1991.
- [DMAMS10] J. Digne, J. M. Morel, N. Audfray, and C. Mehdi-Souzani. The level set tree on meshes. In *Proc. 3DPVT*, 2010.
- [EB] M. M. Bronstein B. Bustos T. Darom R. Horaud I. Hotz Y. Keller J. Keustermans A. Kovnatsky R. Litman J. Reininghaus I. Sipiran D. Smeets P. Suetens D. Vandermeulen A. Zaharescu V. Zobel E. Boyer, A. M. Bronstein. Shrec 2011: robust feature detection and description benchmark. In *EUROGRAPHICS Workshop on 3D Object Retrieval (3DOR)*.
- [EK01] A. Elad and R. Kimmel. Bending invariant representations for surfaces. In *Proc. CVPR*, pages 168–174, 2001.
- [Ela10] M. Elad. *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer, 2010.
- [GBAL09] K. Gebal, J. Andreas Bærentzen, H. Aanæs, and R. Larsen. Shape analysis using the auto diffusion function. *Comput. Graph. Forum*, 28(5):1405–1413, 2009.
- [GF09] A. Golovinskiy and T. Funkhouser. Consistent segmentation of 3d models. *Computers & Graphics*, 33(3):262–269, 2009.
- [HKG11] Q. Huang, V. Koltun, and L. Guibas. Joint shape segmentation with linear programming. In *TOG*, volume 30, page 125, 2011.
- [KLF11] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. *TOG*, 30(4):79, 2011.

- [Kuh55] Harold W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.
- [KZHCO10] O.V. Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. In *Computer Graphics Forum*, volume 20, pages 1–23, 2010.
- [LBB11] R. Litman, A. M. Bronstein, and M. M. Bronstein. Diffusion-geometric maximally stable component detection in deformable shapes. *Computers & Graphics*, 35(3):549 – 560, 2011.
- [Lév06] B. Lévy. Laplace-beltrami eigenfunctions towards an algorithm that understands geometry. In *Proc. SMI*, 2006.
- [LF09] Y. Lipman and T. Funkhouser. Mobius voting for surface correspondence. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 28(3), August 2009.
- [MHK⁺08] Diana Mateus, Radu P. Horaud, David Knossow, Fabio Cuzzolin, and Edmond Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proc. CVPR*, 2008.
- [MS05] Facundo Memoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.*, 5(3):313–347, 2005.
- [NBCW⁺11] A. Nguyen, M. Ben-Chen, K. Welnicka, Y. Ye, and L. Guibas. An optimization approach to improving collections of shape maps. In *Computer Graphics Forum*, volume 30, pages 1481–1491, 2011.
- [Nes07] Y. Nesterov. Gradient methods for minimizing composite objective function. In *CORE Discussion Paper 2007/76, Center for Operations Research and Econometrics (CORE)*. Catholic University of Louvain, Louvain-la-Neuve, Belgium, 2007.
- [OBCS⁺12] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: A flexible representation of maps between shapes. *TOG*, 31(4), 2012.
- [OMMG10] M. Ovsjanikov, Q. Mérigot, F. Mémoli, and L. Guibas. One point isometric matching with the heat kernel. In *Computer Graphics Forum*, volume 29, pages 1555–1564, 2010.
- [PBB11] J. Pokrass, A. M. Bronstein, and M. M. Bronstein. A correspondence-less approach to matching of deformable shapes. In *Proc. SSVM*, 2011.
- [Rus07] R. M. Rustamov. Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proc. of SGP*, pages 225–233, 2007.

- [SBS12] P. Sprechmann, A. M. Bronstein, and G. Sapiro. Learning efficient structured sparse models. In *Proc. ICML*, 2012.
- [SOG09] J. Sun, M. Ovsjanikov, and L. J. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. SGP*, 2009.
- [SY12] Y. Sahillioglu and Y. Yemez. Coarse-to-fine combinatorial matching for dense isometric shape correspondence. In *Computer Graphics Forum*, 2012.
- [TBW⁺11] A. Tevs, A. Berner, M. Wand, I. Ihrke, and H.P. Seidel. Intrinsic shape matching by planned landmark sampling. In *Computer Graphics Forum*, volume 30, pages 543–552, 2011.
- [Tib96] R. Tibshirani. Regression shrinkage and selection via the LASSO. *J. Royal Stat. Society: Series B*, 58(1):267–288, 1996.
- [Tse01] P. Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of optimization theory and applications*, 109(3):475–494, 2001.
- [VKTS⁺11] O. Van Kaick, A. Tagliasacchi, O. Sidi, H. Zhang, D. Cohen-Or, L. Wolf, and G. Hamarneh. Prior knowledge for part correspondence. In *Computer Graphics Forum*, volume 30, pages 553–562, 2011.
- [ZBVH09] A. Zaharescu, E. Boyer, K. Varanasi, and R. Horaud. Surface feature detection and description with applications to mesh matching. In *Proc. CVPR*, 2009.
- [ZWW⁺10] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios. Dense non-rigid surface registration using high-order graph matching. In *Proc. CVPR*, 2010.